
Graph-based Generalization Bounds for Learning Binary Relations

Ben London

University of Maryland
College Park, MD 20742, USA
blondon@cs.umd.edu

Bert Huang

University of Maryland
College Park, MD 20742, USA
bert@cs.umd.edu

Lise Getoor

University of Maryland
College Park, MD 20742, USA
getoor@cs.umd.edu

Abstract

We investigate the generalizability of learned *binary relations*: functions that map pairs of instances to a logical indicator. This problem has application in numerous areas of machine learning, such as ranking, entity resolution and link prediction. Our learning framework incorporates an example labeler that, given a sequence X of n instances and a desired training size m , subsamples m pairs from $X \times X$ without replacement. The challenge in analyzing this learning scenario is that pairwise combinations of random variables are inherently dependent, which prevents us from using traditional learning-theoretic arguments. We present a unified, graph-based analysis, which allows us to analyze this dependence using well-known graph identities. We are then able to bound the generalization error of learned binary relations using Rademacher complexity and algorithmic stability. The rate of uniform convergence is partially determined by the labeler’s subsampling process. We thus examine how various assumptions about subsampling affect generalization; under a natural random subsampling process, our bounds guarantee $\tilde{O}(1/\sqrt{n})$ uniform convergence.

1 Introduction

We investigate the generalizability of a learned *binary relation*: a characteristic function $r : \mathcal{X}^2 \rightarrow \{\pm 1\}$ that indicates whether the input pair satisfies a relation. For example, if r is an equivalence relation,

then $r(x, x') = 1$ indicates that $x \equiv x'$; if r is a total ordering, then $r(x, x') = 1$ means that $x \leq x'$. Binary relations are found in many learning problems, such as ranking (total ordering), entity resolution (equivalence) and link prediction. There has been significant research in each of these fields individually, yet no unified view of the learning problem. We formulate the learning objective (in Section 2) as inductive inference on a product space \mathcal{X}^2 , where instances are drawn from an arbitrary distribution over \mathcal{X} . Given a set of n *independently and identically distributed* (i.i.d.) instances $X \in \mathcal{X}^n$, as well as a subset from X^2 that has been labeled by r , our goal is to produce a hypothesis $h : \mathcal{X}^2 \rightarrow \{\pm 1\}$ that, with high probability, has low error w.r.t. r .

The primary challenge in analyzing this learning setup is reasoning about the dependency structure of a pairwise training set. Even though the instances are i.i.d., the training examples may not be; since each instance may appear in multiple examples, those that involve a common instance are necessarily dependent. This dependence makes the analysis of generalization non-trivial, since it violates the fundamental assumption of independence used in classical statistics and learning theory, rendering existing results incompatible. When the training set does not simply contain all pairs of instances, its dependency structure is difficult to analyze.

In Section 3, we introduce graphical representations of the training set and its dependency structure. The training set is viewed as a graph G , in which each instance is a vertex and each pairwise example is an (undirected) edge. The dependency structure is represented by the corresponding line graph G_D , since edges that share a common vertex (instance) are adjacent in G_D . Casting the dependency structure as a graph allows us to leverage the vast literature of graph theory, which makes our analysis cleaner. The “amount of dependence” can be quantified by the chromatic number of G_D , since the chromatic number is the minimal number of independent sets. Using well-known

chromatic properties, we are able to upper bound this quantity. We then use graph-based arguments in Section 4 to bound the generalization error of learned binary relations, using Rademacher complexity or algorithmic stability. We show that the empirical error converges to the expected error at a rate of $O(\sqrt{\rho/m})$, where m is the number of labeled training pairs and ρ is the maximum frequency of any instance in this set. This ratio depends on how pairs are subsampled (irrespective of their values). Consequently, in Section 5, we explore several learning scenarios in which the subsampling process is working for, against, or is agnostic to the learner (i.e., random). Using a reduction to a random graph model, we prove interesting results for the agnostic scenario. We thus provide a novel analysis of the relationship between subsampling and the rate of uniform convergence for learned binary relations.

1.1 Related Work

Goldman et al. [1993] were the first to address learning a binary relation. Their learning setting is entirely different from ours, in that there is no underlying distribution from which instances are generated, and they are not interested in generalization to unseen data. In their setting, the problem size is polynomial in a finite number of objects. Our learning setting subsumes this one.

Our main analytic tools come from two areas of learning theory: Rademacher complexity and algorithmic stability. The canonical Rademacher bounds are due to Koltchinskii and Panchenko [2002] and Bartlett and Mendelson [2003], while the canonical stability bounds are due to Bousquet and Elisseeff [2002]. Though all of these techniques fundamentally rely on independence, there have been recent applications to non-i.i.d. problems. Inspired by Janson [2004], Usunier et al. [2006] developed a theory of *chromatic* Rademacher analysis, using it to derive generic risk bounds for dependent data, with bipartite ranking as a motivating example. (Ralaivola et al. [2010] used a similar chromatic analysis to derive PAC-Bayes bounds.) Our analysis draws on this work, though we delve deeper into the setting of pairwise prediction. More recently, Mohri and Rostamizadeh [2009, 2010] used an “independent blocking” technique to derive both Rademacher- and stability-based risk bounds for time series data drawn from a strongly-mixing stationary process, though this is an entirely different form of dependence.

A number of authors [Bar-Hillel and Weinshall, 2003, Cl  men  on et al., 2008, Jin et al., 2009] have developed risk bounds for problems involving pairwise prediction (or learning with pairwise constraints), though they all assume that the training set contains all pairwise combinations of the instances, ignoring the more inter-

esting problem of learning from a (sparse) subsample. Our results are not only more general, but they offer insight into how the subsampling process affects the rate of convergence. Agarwal and Niyogi [2009] consider a random subsampling process similar to one we discuss and derive risk bounds for ranking using algorithmic stability. We provide an alternate analysis that is clean, natural and interpretable.

2 Preliminaries

This section introduces the necessary background concepts. To clarify certain probabilities, we let $\Pr_\omega[\cdot]$ denote the probability of an event w.r.t. a random draw of ω from an implicit sample space Ω . Similarly, let $\mathbb{E}_\omega[\cdot]$ denote the expectation of a random variable w.r.t. a random draw of $\omega \in \Omega$, according to an implicit probability distribution.

2.1 Problem Setting

Let \mathcal{X} denote an abstract instance space (e.g., \mathcal{X} could be a subset of Euclidean space). We are interested in learning a binary relation $r : \mathcal{X}^2 \rightarrow \{\pm 1\}$. We will limit our analysis to relations that are *reflexive*, and either *symmetric* or *antisymmetric*, examples of which include equivalence and total ordering.

We define the learning process as follows. We are given a sequence of instances $X \triangleq (x_1, \dots, x_n) \in \mathcal{X}^n$, drawn independently and identically from an arbitrary distribution over \mathcal{X} . We are also given access to a *labeler* S , a black-box process that returns a subset of pairwise labeled examples. More precisely, given an input X and a training size m , $S_m(X)$ returns some subset of X^2 (sampled without replacement) that has been labeled according to r . In practice, this could be a crowd-sourcing application or a targeted surveying process. The subset returned by $S_m(X)$ is sampled according to a process that is independent of the instance values X . Restricting our theory to consider only labelers that determine *which* pairs to label, independent of X , ensures that the observed data is drawn from the original distribution; otherwise, the labeler could introduce bias. Within this restriction, the subsampling process remains unknown. It may be adversarial or benign, aiming to either weaken or improve generalization; it may also be agnostic, selecting pairs according to some random process. We will return to the topic of subsampling in Section 5.

Due to our assumption of (anti)symmetry, only one example per pair is required, since the converse (or contrapositive) can be inferred from context. This means that m is naturally upper bounded by $\binom{n}{2}$. Upon receiving a labeled dataset $Z \leftarrow S_m(X)$, we invoke a

learning algorithm A to obtain an *inductive* hypothesis h from a class \mathcal{H} of functions from \mathcal{X}^2 to $\{\pm 1\}$. We let h_Z (or, more explicitly, $h_Z \leftarrow A(Z)$) denote the hypothesis trained on the dataset Z .

2.2 Generalization

For a given *cost function* $c : \mathbb{R}^2 \rightarrow \mathbb{R}^+$, define the *loss* ℓ of a hypothesis h w.r.t. a pair $z \triangleq (x, x')$ as $\ell(h, z) \triangleq c(r(z), h(z))$. Though there are various cost functions, for learning binary relations, we are most interested in the so-called *0-1 loss*, $\ell_1(h, z) \triangleq \mathbb{1}[r(z) \neq h(z)]$.

The quantity we are primarily concerned with is the *generalization error*, or *risk*. This is the expected loss of a learned hypothesis h_Z w.r.t. a random pair $\tilde{z} \triangleq (x, x')$, where x and x' are sampled independently and identically from the underlying distribution over \mathcal{X} . We denote this quantity by $R(h_Z) \triangleq \mathbb{E}_{\tilde{z}}[\ell(h_Z, \tilde{z})]$. In practice, we can estimate the true risk with the *training error*, computed as the average loss over all examples in the training set. We denote this by $\hat{R}(h_Z) \triangleq \frac{1}{m} \sum_{z \in Z} \ell(h_Z, z)$.

Given an empirical risk estimate, we would like to bound its deviation from the true risk. We will refer to this quantity as the *defect*, which we denote by $D(h_Z) \triangleq R(h_Z) - \hat{R}(h_Z)$. We will use the canonical *probably approximately correct* (PAC) framework [Valiant, 1984], in which we allow the defect to exceed an arbitrary $\epsilon > 0$ with probability $\delta \in (0, 1)$. There are a number of ways of analyzing this probability, of which we will focus on hypothesis complexity and algorithmic stability. Analyses of this nature typically aim to bound

$$\Pr \left[\sup_{h \in \mathcal{H}} D(h) \geq \epsilon \right] \quad \text{or} \quad \Pr \left[\sup_{h_Z \leftarrow A_Z} D(h_Z) \geq \epsilon \right].$$

Solving for ϵ , one obtains a probabilistic upper bound on the true risk, parameterized by δ .

3 Graphical Representation

In this section, we analyze the learning problem and its inherent dependencies using a graphical representation.

3.1 Training Data

Recall that the training set, returned by the labeler S , is an arbitrary subset of X^2 that has been labeled according to r . We can represent this as a graph $G \triangleq (V, E)$, in which each vertex $v_i \in V$ corresponds to an instance x_i , and each training example $(x_i, x_j, r(x_i, x_j)) \in Z$ defines an (un)directed edge $(i, j) \in E$. Thus, the subsampling pattern reduces to

a graph on the instances X . Note that, if r is antisymmetric, then G will be directed; however, since only one example per pair is needed, the corresponding undirected graph is simple (i.e., not a multigraph). This graph representation simplifies the analysis of generalization and allows the labeler's subsampling to be viewed as one of many well-studied graph generation processes.

For example, we can consider an agnostic, random labeler that selects pairs uniformly at random. Using the above graphical representation, this subsampling process can be modeled by the Erdős-Rényi random graph model $\mathcal{G}(n, m)$. In this model, a graph with n vertices and m edges is chosen uniformly from the set of all such graphs. (This model differs slightly from the popular $\mathcal{G}(n, p)$ model, in which each edge is realized independently with probability p .) We analyze generalization using various labeler scenarios, including the $\mathcal{G}(n, m)$ labeler, in Section 5.

3.2 Dependency Structure

For each instance x_i , define a random variable X_i , and recall that these are i.i.d. For each example pair (x_i, x_j) found in Z , define a random variable $Z_{i,j} \triangleq (X_i, X_j)$. Because each instance may appear in multiple pairs, we have that these random variables are not mutually independent. To make this more concrete, consider the set $\{Z_{1,2}, Z_{2,3}, Z_{3,4}\}$. Clearly, $Z_{1,2}$ and $Z_{2,3}$ are dependent, since they include a common variable, X_2 . Now consider variables $Z_{1,2}$ and $Z_{3,4}$, and note that $\{1, 2\} \cap \{3, 4\} = \emptyset$. Observing (X_1, X_2) reveals nothing about (X_3, X_4) ; thus, $\mathbb{P}(Z_{3,4} \mid Z_{1,2}) = \mathbb{P}(Z_{3,4})$, and vice versa, so they are mutually independent.

We will represent the dependency structure using a graphical representation due to Erdős and Lovász [1975], known as a *dependency graph*.

Definition 1 (Dependency Graph). Let $\mathbf{Z} \triangleq \{Z_i\}_{i=1}^n$ be a set of random variables with joint distribution $\mathbb{P}(\mathbf{Z})$, and let $G_D \triangleq (V, E)$ be a graph, with $V \triangleq [n]$. Then G_D forms a *dependency graph* w.r.t. \mathbf{Z} if every *independent set* (i.e., set of non-adjacent vertices) $I \subseteq V$ satisfies $\mathbb{P}(\{Z_i\}_{i \in I}) = \prod_{i \in I} \mathbb{P}(Z_i)$.

In the context of learning binary relations, we construct a dependency graph G_D in which each vertex $v_{i,j}$ is connected to vertices $\{v_{k,\ell} : (k=i) \vee (\ell=j)\}$ —i.e., the set of vertices that involve instances x_i or x_j . As a result, any independent set of vertices (in the graph-theoretic sense) is a set of independent random variables. To better understand the structure of this dependency graph, it helps to recall the graph G defined in the previous section; G_D is in fact its corresponding *line graph*, a graph representing the adjacencies

cies between edges.¹ Therefore, every independent set in G_D is a *matching* in G .

3.3 Chromatic Properties

In this section, we discuss the chromatic properties of the above graphical representations, which will introduce a key lemma used in our generalization bounds. For completeness, we first review some background on graph coloring. For the following definitions, let $G \triangleq (V, E)$ be an arbitrary undirected graph.

Definition 2 (Vertex Coloring). A *proper k -vertex-coloring* (often simply referred to as simply a k -coloring) is a mapping from V to a set of k color classes, such that no two adjacent vertices have the same color. Equivalently, it is a partitioning $\{C_j : C_j \subseteq V\}_{j=1}^k$, such that $\bigcup_{j=1}^k C_j = V$, $\bigcap_{j=1}^k C_j = \emptyset$, and every subset C_j is independent. The *chromatic number* $\chi(G)$ is the minimum number of colors needed to achieve a proper coloring.

Definition 3 (Edge Coloring). A *proper k -edge-coloring* is a mapping from E to a set of k color classes, such that no two coincident edges have the same color. The *chromatic index* $\chi'(G)$ is the minimum number of colors needed to achieve a proper edge coloring.

Theorem 1 (Vizing, 1964). *If G has maximum degree $\Delta(G)$, then $\Delta(G) \leq \chi'(G) \leq \Delta(G) + 1$.*

For a dependency graph, the chromatic number can be viewed roughly as the “amount of dependence”. The lower the chromatic number, the more independence. If the variables are i.i.d., then the chromatic number is 1. Coloring the vertices of the dependency graph G_D described in Section 3.2 can be reduced to coloring the edges of the graph G described in Section 3.1, since one is simply the line graph of the other. The chromatic index of G is equal to the chromatic number of G_D , so bounding one quantity bounds the other. Although there are many graphs for which $\chi'(G) = \Delta(G)$, determining the chromatic index of an arbitrary graph is NP-hard, so we will rely on the upper bound. We can therefore state the amount of dependence in terms of the chromatic index of G . Note that the maximum degree $\Delta(G)$ is simply the maximum frequency of any instance in the training set, which yields the following lemma.

Lemma 1. *Let $X \in \mathcal{X}^n$ be a set of n i.i.d. instances, and $Z \leftarrow S_m(X)$ an arbitrary training set of size m , with maximum instance frequency ρ . Let G_D be the corresponding dependency graph of Z . Then, $\chi(G_D) \leq \rho + 1$.*

¹If G is directed, then G_D is the line graph of its corresponding undirected graph.

4 Generalization Bounds

In this section, we develop risk bounds for learning binary relations using both the Rademacher complexity and algorithmic stability.

4.1 Concentration Inequalities

A key component in any generalization analysis is the concentration of random variables. We now present two tail bounds that will be used in our proofs.

Theorem 2 (McDiarmid, 1989). *Let $\mathbf{Z} \triangleq \{Z_i\}_{i=1}^n$ be a set of i.i.d. random variables that take values in \mathcal{Z} . Let $f : \mathcal{Z}^n \rightarrow \mathbb{R}$ be a measurable function for which there exist constants $\{\alpha_i\}_{i=1}^n$ such that, for any $i \in [n]$, and any inputs $Z, Z' \in \mathcal{Z}^n$ that differ only in the i^{th} variable, $|f(Z) - f(Z')| \leq \alpha_i$. Then, for any $\epsilon > 0$,*

$$\Pr[f(\mathbf{Z}) - \mathbb{E}[f(\mathbf{Z})] \geq \epsilon] \leq \exp\left(\frac{-2\epsilon^2}{\sum_{i=1}^n \alpha_i^2}\right).$$

Theorem 3 (Usunier et al., 2006). *Let G_D be a dependency graph (Definition 1) for a set of random variables $\mathbf{Z} \triangleq \{Z_i\}_{i=1}^n$ that take values in \mathcal{Z} . Let $\{\mathbf{Z}_j\}_{j=1}^{\chi(G_D)}$ denote the subsets induced by an optimal proper coloring of G_D , and let $n_j \triangleq |\mathbf{Z}_j|$. Finally, let $f : \mathcal{Z}^n \rightarrow \mathbb{R}$ be a measurable function where: (a) there exist functions $\{f_j : \mathcal{Z}^{n_j} \rightarrow \mathbb{R}\}_{j=1}^{\chi(G_D)}$, such that $f(\mathbf{Z}) = \sum_j f_j(\mathbf{Z}_j)$; (b) there exists a constant α such that every f_j is α -Lipschitz w.r.t. the Hamming metric. Then, for any $\epsilon > 0$,*

$$\Pr[f(\mathbf{Z}) - \mathbb{E}[f(\mathbf{Z})] \geq \epsilon] \leq \exp\left(\frac{-2\epsilon^2}{\chi(G_D)n\alpha^2}\right).$$

4.2 Rademacher Complexity

Informally, the Rademacher complexity measures a hypothesis class’ expressive power, quantified by its ability to fit a random signal. We slightly adapt the traditional definition to better suit our learning context.

Definition 4 (Rademacher Complexity). Let \mathcal{X} be an instance space, and \mathcal{Z} some alternate space. Let $\varphi : \mathcal{X}^n \rightarrow \mathcal{Z}^m$ denote a mapping from n instances from \mathcal{X} to m instances from \mathcal{Z} . Let σ be a set of *Rademacher variables* that independently take values in $\{\pm 1\}$ with equal probability. For a class \mathcal{F} of functions from \mathcal{Z} to \mathbb{R} , define the *empirical Rademacher complexity* of $\mathcal{F} \circ \varphi$, w.r.t. instances $X \in \mathcal{X}^n$, as

$$\mathfrak{R}_X(\mathcal{F} \circ \varphi) \triangleq \mathbb{E}_\sigma \left[\frac{2}{m} \sup_{f \in \mathcal{F}} \left| \sum_{i=1}^m \sigma_i f(z_i) \right| \right].$$

Define the *Rademacher complexity* of $\mathcal{F} \circ \varphi$, w.r.t. i.i.d. samples of size n , as $\mathfrak{R}_n(\mathcal{F} \circ \varphi) \triangleq \mathbb{E}_X[\mathfrak{R}_X(\mathcal{F} \circ \varphi)]$.

In our scenario, the data may exhibit dependencies; yet, as we will see, this does not affect the traditional symmetry argument used in Rademacher analysis.

Theorem 4. *Let $X \in \mathcal{X}^n$ be an i.i.d. sample of n instances, and $Z \leftarrow S_m(X)$ an arbitrary training set of size m , with maximum instance frequency ρ . Let \mathcal{F} be a class of functions from $\mathcal{Z} \triangleq \mathcal{X}^2$ to $[0, 1]$. Then, for any $n \geq 2$, any $m \geq 1$, any $f \in \mathcal{F}$, and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over draws of X ,*

$$\mathbb{E}_{\tilde{z}}[f(\tilde{z})] \leq \frac{1}{m} \sum_{z \in Z} f(z) + \mathfrak{R}_n(\mathcal{F} \circ S_m) + \sqrt{\frac{\rho + 1}{2m} \ln \frac{1}{\delta}}. \quad (1)$$

Proof With G_D as the dependency graph of Z , we invoke an optimal proper coloring, which partitions Z into $\chi(G_D) \leq \rho + 1$ (via Lemma 1) independent sets. We can consequently express the defect $D(f) \triangleq \mathbb{E}_{\tilde{z}}[f(\tilde{z})] - \frac{1}{m} \sum_{z \in Z} f(z)$ as a sum over functions of independent random variables. It is easy to show that each of these functions is $(1/m)$ -Lipschitz w.r.t. the Hamming metric. We therefore apply Theorem 3 and obtain

$$\Pr_X \left[\sup_{f \in \mathcal{F}} D(f) - \mathbb{E}_X[\sup_{f \in \mathcal{F}} D(f)] \geq \epsilon \right] \leq \exp \left(\frac{-2\epsilon^2 m}{\rho + 1} \right), \quad (2)$$

To bound $\mathbb{E}_X[\sup_{f \in \mathcal{F}} D(f)]$, we start by imagining a “ghost sample” X' of n i.i.d. instances that have been labeled using the same pattern as Z to create Z' . Note that $\mathbb{E}_{\tilde{z}}[f(\tilde{z})] = \mathbb{E}_{X'}[\frac{1}{m} \sum_{z' \in Z'} f(z')]$ via linearity of expectation, since all z' have the same marginal distribution. Using the ghost sample and Jensen’s inequality, we have that

$$\begin{aligned} \mathbb{E}_X[\sup_{f \in \mathcal{F}} D(f)] &= \mathbb{E}_X \left[\sup_{f \in \mathcal{F}} \mathbb{E}_{\tilde{z}}[f(\tilde{z})] - \frac{1}{m} \sum_{z \in Z} f(z) \right] \\ &= \mathbb{E}_X \left[\sup_{f \in \mathcal{F}} \mathbb{E}_{X'} \left[\frac{1}{m} \sum_{z' \in Z'} f(z') \right] - \frac{1}{m} \sum_{z \in Z} f(z) \right] \\ &\leq \mathbb{E}_{X, X'} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m f(z'_i) - f(z_i) \right]. \end{aligned} \quad (3)$$

Now define a set of Rademacher variables σ , and define random variables $\{Z_i\}_{i=1}^m$ and $\{Z'_i\}_{i=1}^m$ for Z and Z' sets respectively. Because they are labeled using the same pattern, and have isomorphic dependency graphs, we have that $\mathbb{P}(Z_1, \dots, Z_m) = \mathbb{P}(Z'_1, \dots, Z'_m)$. In fact, if we exchange any Z_i and Z'_i , we have that $\mathbb{P}(Z_1, \dots, Z'_i, \dots, Z_m) = \mathbb{P}(Z'_1, \dots, Z_i, \dots, Z'_m)$. Thus, since every draw of σ occurs with equal probability, we

have via symmetry that

$$\begin{aligned} \text{Eq. (3)} &\leq \mathbb{E}_{X, X', \sigma} \left[\sup_{f \in \mathcal{F}} \frac{1}{m} \sum_{i=1}^m \sigma_i (f(z'_i) - f(z_i)) \right] \\ &\leq \mathbb{E}_{X, \sigma} \left[\sup_{f \in \mathcal{F}} \frac{2}{m} \left| \sum_{i=1}^m \sigma_i f(z'_i) \right| \right] \leq \mathfrak{R}_n(\mathcal{F} \circ S_m). \end{aligned}$$

and so $\mathbb{E}_X[\sup_{f \in \mathcal{F}} D(f)] \leq \mathfrak{R}_n(\mathcal{F} \circ S_m)$. To obtain Equation 1, we simply set Equation 2 equal to δ and solve for ϵ . ■

It is possible to derive a similar risk bound for the empirical Rademacher complexity by applying Theorem 3 to the difference of $\mathfrak{R}_n(\mathcal{F} \circ S_m) - \mathfrak{R}_X(\mathcal{F} \circ S_m)$. We omit this proof to save space, since the remainder of our work does not require the empirical Rademacher complexity.

To make these bounds more functional, we will replace \mathfrak{R}_n with an empirically verifiable quantity. For certain hypothesis classes, it is possible to show that the Rademacher complexity is bounded by a function of the model parameters. One such class of hypotheses is *reproducing kernel Hilbert spaces* (RKHS), which subsume the popular support vector machine (SVM) [Cristianini and Shawe-Taylor, 2000]. Formally, for some mapping $\phi : \mathcal{Z} \rightarrow \tilde{\mathcal{Z}}$, where \mathcal{Z} is an instance space and $\tilde{\mathcal{Z}}$ is a Hilbert space, endowed with an inner product $\langle \cdot, \cdot \rangle$, a kernel $\kappa : \mathcal{Z}^2 \rightarrow \mathbb{R}$ is a function such that, for all $(z, z') \in \mathcal{Z}^2$, $\kappa(z, z') = \langle \phi(z), \phi(z') \rangle$. The only requirement is that the kernel’s Gram matrix $\mathbf{K} : \mathbf{K}_{i,j} = \kappa(z_i, z_j)$ be symmetric and positive semidefinite. RKHS hypotheses are generally of the form $h_{\kappa, Z}(z) \triangleq \sum_{z' \in Z}^m \kappa(z, z')$, where $Z \in \mathcal{Z}^m$ is a set of reference points from the problem domain (e.g., support vectors). We may reasonably assume that the norm of the kernel mapping is uniformly bounded by a constant C ; i.e., the mapped data is contained within a ball of radius C . We denote the class of kernel functions by \mathcal{H}_κ . Borrowing results from Koltchinskii and Panchenko [2002] and Bartlett and Mendelson [2003], we are able to prove the following risk bounds for kernel-based hypotheses with bounded kernels, in the context of learning binary relations.

For the following, we use the *ramp loss* as a surrogate for the 0-1 loss. For a given $\gamma > 0$, a real-valued hypothesis $h : \mathcal{X}^2 \rightarrow \mathbb{R}$ and an example z , define the ramp loss as $\ell_\gamma(h, z) \triangleq \min\{\max\{0, 1 - r(z)h(z)/\gamma\}, 1\}$. To differentiate this from the 0-1 loss when dealing with risk metrics, we henceforth use a superscript $\mathbb{1}$ or γ .

Theorem 5. *Let $X \in \mathcal{X}^n$, $Z \leftarrow S_m(X)$ and ρ be as defined in Theorem 4. Let $h_{\kappa, Z}$ be a RKHS hypothesis trained on Z , such that $\sup_z \|\phi(z)\| \leq C$. Then, for any $n \geq 2$, any $m \geq 1$, any $\gamma > 0$, and any $\delta \in (0, 1)$,*

with probability at least $1 - \delta$ over draws of X ,

$$R^1(h_{\kappa,Z}) \leq \hat{R}^\gamma(h_{\kappa,Z}) + \frac{4C}{\gamma\sqrt{m}} + \sqrt{\frac{\rho+1}{2m} \ln \frac{1}{\delta}}. \quad (4)$$

Proof Note that ℓ_γ dominates ℓ_1 , and thus $R^1(h_{\kappa,Z}) \leq R^\gamma(h_{\kappa,Z})$. Since the ramp loss is bounded by $[0, 1]$, we apply Theorem 4 and obtain

$$R^1(h_Z) \leq \hat{R}^\gamma(h_Z) + \mathfrak{R}_n(\ell_\gamma \circ \mathcal{H}_\kappa \circ S_m) + \sqrt{\frac{\rho+1}{2m} \ln \frac{1}{\delta}}.$$

To bound $\mathfrak{R}_n(\ell_\gamma \circ \mathcal{H}_\kappa \circ S_m)$, we use Talagrand's contraction lemma [Ledoux and Talagrand, 1991] and borrow a result from Bartlett and Mendelson [2003, Lemma 22], from which we obtain

$$\mathfrak{R}_n(\ell_\gamma \circ \mathcal{H}_\kappa \circ S_m) \leq \frac{2}{\gamma} \mathfrak{R}_n(\mathcal{H}_\kappa \circ S_m) \leq \frac{2}{\gamma} \cdot \frac{2\sqrt{\text{tr}(\mathbf{K})}}{m}.$$

Using Cauchy-Schwarz, we can bound the trace of the kernel's Gram matrix as

$$\text{tr}(\mathbf{K}) = \sum_{z \in Z} \langle \phi(z), \phi(z) \rangle \leq \sum_{z \in Z} \|\phi(z)\|^2 \leq mC^2.$$

We therefore have that $\mathfrak{R}_n(\ell_\gamma \circ \mathcal{H}_\kappa \circ S_m) \leq \frac{4C}{\gamma\sqrt{m}}$. ■

Note that this analysis slightly improves upon that of Usunier et al. [2006] in that we use the regular Rademacher complexity instead of their so-called *fractional Rademacher complexity*. Because of this, our Rademacher term is $O(1/\sqrt{m})$, compared to $O(\sqrt{\rho/m})$ using the fractional Rademacher complexity; note that $1/\sqrt{m} \leq \sqrt{\rho/m}$ for all $\rho \geq 1$.

4.3 Algorithmic Stability

This section derives a different generalization bound for learning pairwise relations using our previous graph representations and *algorithmic stability*.

Definition 5 (Uniform Stability). For a training set Z , let Z' be a duplicate of Z with the i^{th} example removed. A learning algorithm A has *uniform stability* β w.r.t. a loss function ℓ if, for any $Z \in \mathcal{Z}^m$, and any $i \in [m]$, A returns hypotheses $h_Z \leftarrow A(Z)$ and $h_{Z'} \leftarrow A(Z')$ such that

$$\sup_{\tilde{z} \in \mathcal{Z}} |\ell(h_Z, \tilde{z}) - \ell(h_{Z'}, \tilde{z})| \leq \beta.$$

In other words, excluding any single example from training will increase the loss, w.r.t. any test example \tilde{z} , by at most β . Of course, β must be a function of the size of the training set; indeed, we will later show that generalization is only possible when $\beta = O(1/m)$. To highlight this dependence, we henceforth use the notation β_m . Using this notion of stability, we now derive alternate risk bounds for learning binary relations.

Lemma 2 (Bousquet and Elisseeff, 2002). Let A be a learning algorithm with uniform stability β w.r.t. a loss function ℓ , where ℓ is upper bounded by M . Then, for any $i \in [m]$, and any training sets $Z, Z' \in \mathcal{Z}^m$ that differ only in the value of the i^{th} example, A returns hypotheses $h_Z \leftarrow A(Z)$ and $h_{Z'} \leftarrow A(Z', G)$ such that

$$|D(h_Z) - D(h_{Z'})| \leq 4\beta_m + \frac{M}{m}.$$

Theorem 6. Let $X \in \mathcal{X}^n$ be an i.i.d. sample of n instances, and $Z \leftarrow S_m(X)$ an arbitrary training set of size m , with maximum instance frequency ρ . Let A be a learning algorithm with uniform stability β w.r.t. a loss function ℓ (upper bounded by M), and let $h_Z \leftarrow A(Z)$. Then, for any $n \geq 2$, any $m \geq 1$, and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over draws of X ,

$$R(h_Z) \leq \hat{R}(h_Z) + 4\rho\beta_m + (4m\beta_m + M)\sqrt{\frac{\rho}{m} \ln \frac{1}{\delta}}. \quad (5)$$

Proof We begin by showing that the defect satisfies the conditions of McDiarmid's inequality (Theorem 2). By Lemma 2, replacing any single *example* will change the defect by at most $4\beta_m + M/m$. However, if we replace any single *instance* x_i , this will affect up to ρ_i examples, where ρ_i denotes the frequency of x_i in the training set. We therefore have that replacing any x_i has Lipschitz constant $\alpha_i = \rho_i(4\beta_m + M/m)$. Since the instances are i.i.d., we can apply McDiarmid's inequality and obtain

$$\begin{aligned} \Pr_X [D(h_Z) - \mathbb{E}_X[D(h_Z)]] &\leq \exp\left(\frac{-2\epsilon^2}{\sum_{i=1}^n \rho_i^2 (4\beta_m + M/m)^2}\right) \\ &\leq \exp\left(\frac{-2\epsilon^2 m^2}{(4m\beta_m + M)^2 \rho \sum_{i=1}^n \rho_i}\right) \\ &= \exp\left(\frac{-\epsilon^2 m}{(4m\beta_m + M)^2 \rho}\right). \end{aligned}$$

The last line is due to the *handshaking lemma*, which states that the sum of the degrees (i.e., instance frequencies) in a graph is equal to twice the number of edges (i.e., examples). Setting the above equal to δ and solving for ϵ , we have that

$$R(h_Z) \leq \hat{R}(h_Z) + \mathbb{E}_X[D(h_Z)] + (4m\beta_m + M)\sqrt{\frac{\rho}{m} \ln \frac{1}{\delta}}.$$

What remains is to upper bound the expected defect. Using linearity of expectation, we can state this as

$$\mathbb{E}_X[D(h_Z)] = \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{X, \tilde{z}}[\ell(h_Z, \tilde{z}) - \ell(h_Z, z_i)].$$

Note that example $z_i = (x, x')$ depends on any examples that share either x or x' , i.e., its neighborhood $\mathcal{N}(z_i)$ in the dependency graph. However, by removing z_i and $\mathcal{N}(z_i)$ from the training set, z_i becomes independent of any of the remaining examples. Accordingly, let $Z' \triangleq Z \setminus \{z_i, \mathcal{N}(z_i)\}$, and let $h_{Z'}$ be the resulting hypothesis. Because we have removed $\deg(z_i) + 1$ examples from training, by uniform stability, we pay a penalty of at most $\beta_m(\deg(z_i) + 1)$ loss per prediction. Further, recall that G_D is the line graph of the graph G defined in Section 3.1. It is therefore easy to show that any node in G_D (i.e., edge in G) has degree equal to the sum of the degrees of its endpoints in G , minus two; hence, $\deg(z_i) + 1 = \deg(x) + \deg(x') - 1 \leq 2\rho - 1$. We therefore have that

$$\begin{aligned} \mathbb{E}_X[D(h_Z)] &\leq \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{X, \tilde{z}}[\ell(h_{Z'}, \tilde{z}) - \ell(h_{Z'}, z_i)] \\ &\quad + 2\beta_m(\deg(z_i) + 1) \\ &= \frac{2\beta_m}{m} \sum_{i=1}^m 0 + \deg(z_i) + 1 \\ &\leq 2\beta_m(2\rho - 1) \leq 4\rho\beta_m, \end{aligned}$$

where the second line follows from symmetry, since z and z_i are now i.i.d. variables. ■

To obtain a non-vacuous bound, we require that $\beta_m = O(1/m)$. This precludes stability w.r.t. the 0-1 loss ℓ_1 , since any algorithm will have either $\beta = 0$ or $\beta = 1$, regardless of m . We therefore use the ramp loss ℓ_γ again. To use the ramp loss, we introduce the notion of *classification stability*. For the following, we consider learning algorithms that output a real-valued hypothesis $h : \mathcal{Z} \rightarrow \mathbb{R}$, where $\text{sgn}(h(z))$ is the predicted label of z .

Definition 6 (Classification Stability). Let Z and Z' be as defined in Definition 5. A learning algorithm A has *classification stability* β if, for any $Z \in \mathcal{Z}^m$, and any $i \in [m]$, A returns real-valued hypotheses $h_Z \leftarrow A(Z)$ and $h_{Z'} \leftarrow A(Z')$ such that

$$\sup_{\tilde{z} \in \mathcal{Z}} |h_Z(\tilde{z}) - h_{Z'}(\tilde{z})| \leq \beta.$$

Lemma 3 (Bousquet and Elisseeff, 2002). *A learning algorithm with classification stability β has uniform stability β/γ w.r.t. the ramp loss ℓ_γ .*

Theorem 7. *Let $X \in \mathcal{X}^n$, $Z \leftarrow S_m(X)$ and ρ be as defined in Theorem 6. Let A be a learning algorithm with classification stability β , and let $h_Z \leftarrow A(Z)$. Then, for any $n \geq 2$, any $m \geq 1$, any $\gamma > 0$, and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over draws of X ,*

$$R^1(h_Z) \leq \hat{R}^\gamma(h_Z) + \frac{4\rho\beta_m}{\gamma} + \left(\frac{4m\beta_m}{\gamma} + 1\right) \sqrt{\frac{\rho}{m} \ln \frac{1}{\delta}}. \quad (6)$$

Proof The proof follows directly from Theorem 6 and Lemma 3, with the ramp loss obviously upper bounded by $M = 1$. ■

The application of this bound still depends on a stability parameter, which is unique to the learning algorithm. As a demonstration, we will focus on the class of kernel methods described in Section 4.2; specifically, SVM classification. Recall that, using stability analysis, generalization is made possible by properties of the learning algorithm, not the complexity of the hypothesis class. For the class of kernel methods in particular, this mechanism is *regularization*. We define a kernel-based regularization algorithm as one of the form

$$A_\kappa(Z) \triangleq \underset{h \in \mathcal{H}_\kappa}{\operatorname{argmin}} \frac{1}{m} \sum_{z \in Z} \ell(h, z) + \lambda \|h\|^2,$$

where $\lambda > 0$ is a regularization parameter. The loss function varies, depending on the application and algorithm. In SVM classification, it is common to minimize the *hinge loss*, defined as $\ell_h(h, z) \triangleq \max\{0, 1 - r(z)h(z)\}$. Denote the empirical hinge risk by \hat{R}^h . Using a stability result from Bousquet and Elisseeff [2002], we obtain a risk bound for SVM classification.

Lemma 4 (Bousquet and Elisseeff, 2002). *An SVM learning algorithm, with $\sup_z \|\phi(z)\| \leq C$ and regularization parameter $\lambda > 0$, has classification stability $\beta_m \leq C^2/(2\lambda m)$.*

Theorem 8. *Let $X \in \mathcal{X}^n$, $Z \leftarrow S_m(X)$ and ρ be as defined in Theorem 6. Let A_κ be an SVM learning algorithm, with $\sup_z \|\phi(z)\| \leq C$ and $\lambda > 0$, and let $h_{\kappa, Z} \leftarrow A_\kappa(Z)$. Then, for any $n \geq 2$, any $m \geq 1$, and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over draws of X ,*

$$R^1(h_{\kappa, Z}) \leq \hat{R}^h(h_{\kappa, Z}) + \frac{2\rho C^2}{\lambda m} + \left(\frac{2C^2}{\lambda} + 1\right) \sqrt{\frac{\rho}{m} \ln \frac{1}{\delta}}. \quad (7)$$

Proof Clearly, for $\gamma = 1$, \hat{R}^h dominates \hat{R}^γ . We therefore apply Theorem 7, with $\gamma = 1$ and $\beta = C^2/(2\lambda m)$. ■

4.4 Discussion of Bounds

Our bounds are dominated by the term ρ/m , where ρ is the maximum instance frequency (equivalently, the maximum degree in G) and m is the number of examples (edges). We refer to the inverse of this ratio as the *effective training size*. Letting ρ_i denote the frequency (i.e., degree) of instance x_i , we have that

$$\frac{\rho}{m} = \frac{2\rho}{\sum_i^n \rho_i} = \frac{2}{\sum_i^n \rho_i / \rho},$$

It is straightforward to show that this quantity is minimized when G is regular—that is, when ρ_i is uniform. In fact, for any regular G , we have that $\rho/m = 2/n$.

Assuming one cannot acquire new examples, one can discard examples to make a regular graph, which gives the optimal ratio. This is equivalent to finding a k -regular spanning subgraph (i.e., k -factor), for some $k \geq 1$. This is not always possible without reducing the number of instances (i.e., vertices), as some graphs do not admit such a k -factor. For example, if the highest degree vertex is adjacent to multiple degree-1 vertices (as in a star graph), then certain vertices will be “isolated” when edges are removed. In fact, the *effective* number of instances n' is the order of the largest k -regular induced subgraph, for some $k \geq 1$; we therefore have that the effective training size is upper bounded by $n'/2$. That said, identifying n' for an arbitrary graph is NP-hard.

It is tempting to think that, by discarding examples to induce a 1-regular subgraph, one can reduce our learning setup to the i.i.d. scenario and consequently apply classical analysis. However, there may be a regular (sub)graph of higher degree that yields a better effective training size. For instance, consider a graph consisting of t disjoint triangles (i.e., $n = 3t$); this graph is already 2-regular, so without pruning edges it has an effective training size $n/2 = 3t/2$; if pruned to a 1-regular subgraph, the effective training size would be just t . Moreover, while the above shows that discarding examples might minimize our bounds, without intimate knowledge of the learning algorithm, hypothesis class or distribution, our bounds may be overly pessimistic in certain scenarios. Stronger assumptions may lead to tighter risk bounds, to support the intuition that more training data—albeit dependent data—will always improve generalization.

5 On Subsampling and the Rate of Uniform Convergence

We have shown that the empirical risk converges to the true risk at a rate of $O(\sqrt{\rho/m})$, depending primarily on the size of the training set and the maximum frequency of any instance. While m may be determined by one’s annotation or computation budget, ρ depends on the subsampling used to select the training set. In this section, we examine the relationship between the subsampling process used by the labeler and the rate of uniform convergence.

Recall that the labeler *cannot* subsample based on the values of the input data, but it can subsample patterns that help or hurt generalization. If the labeler is working *against* the learner, it can select pairs such that

one instance appears in all training examples, meaning $\rho = m$. In this scenario, our bounds indicate that a hypothesis learned from this training set might not generalize. In contrast, if the labeler is working *with* the learner, it can subsample pairs so as to induce a regular label graph, as discussed in the previous section. This would yield an optimal convergence rate of $O(1/\sqrt{n})$, comparable to classical results.

We may also consider a setting in which the subsampling is a random process. For example, if the labeler selects pairs uniformly at random without replacement, then, as previously noted, this process can be modeled by the Erdős-Rényi random graph model. We then have that the rate of convergence is a function of the maximum degree of $\mathcal{G}(n, m)$, since this is equivalent to the maximum instance frequency.

Lemma 5. *Let $G \triangleq (V, E)$ be a graph in $\mathcal{G}(n, m)$, for a given n and m . Then, with probability at least $1 - \delta$, its maximum degree $\Delta(G)$ is upper bounded as*

$$\Delta(G) \leq \frac{2m}{n} \left(1 + \sqrt{\frac{3n}{2m} \ln \frac{n}{\delta}} \right). \quad (8)$$

We provide the proof in Appendix A.

Using this as an upper bound for ρ (since $\rho = \Delta(G)$), we obtain the following corollary of Theorem 5. A similar result can be shown for Theorem 8.

Theorem 9. *Let $X \in \mathcal{X}^n$, $Z \leftarrow S_m(X)$, $h_{\kappa, Z}$, C and λ be as previously defined. If S samples m examples uniformly at random from X^2 , then, for any $n \geq 2$, any $m \geq n/2$, any $\gamma > 0$, and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over draws of X , and $\eta \triangleq 1 + \sqrt{\frac{3n}{2m} \ln \frac{2n}{\delta}}$,*

$$R^1(h_{\kappa, Z}) \leq \hat{R}^\gamma(h_{\kappa, Z}) + \frac{\sqrt{32}C}{\gamma\sqrt{n}} + \sqrt{\frac{\eta+1}{n} \ln \frac{2}{\delta}}. \quad (9)$$

Proof Equation 9 follows from Theorem 5 and Lemma 5, by allowing failure probability $\delta/2$ to Equation 4 and $\delta/2$ to Equation 8. We then simplify the bound by leveraging the fact that $1/m \leq 2/n$. ■

We point out that these bounds have a natural interpretation. Whereas Agarwal and Niyogi [2009] invoke parameterized families of edge distributions, we consider a simple, intuitive learning setup in which the only parameter is the size m of the training set. If $m \geq n \ln n$, then $\eta = O(1)$, and we obtain a uniform convergence rate of $O(1/\sqrt{n})$. For $m \geq n/2$, we have that $\eta = O(\sqrt{\ln n})$, and the rate is still $\tilde{O}(1/\sqrt{n})$.

References

- S. Agarwal and P. Niyogi. Generalization bounds for ranking algorithms via algorithmic stability. *Journal of Machine Learning Research*, 10:441–474, 2009.
- A. Bar-Hillel and D. Weinshall. Learning with equivalence constraints, and the relation to multiclass learning. In *Proceedings of the Conference on Computational Learning Theory*, 2003.
- P. Bartlett and S. Mendelson. Rademacher and gaussian complexities: risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, March 2003.
- O. Bousquet and A. Elisseeff. Stability and generalization. *Journal of Machine Learning Research*, 2: 499–526, March 2002.
- S. Cl  men  on, G. Lugosi, and N. Vayatis. Ranking and empirical minimization of U-statistics. *The Annals of Statistics*, 36(2):844–874, April 2008. ISSN 0090-5364.
- N. Cristianini and J. Shawe-Taylor. *An introduction to support vector machines and other kernel-based learning methods*. Cambridge University Press, March 2000.
- P. Erd  s and L. Lov  sz. Problems and results on 3-chromatic hypergraphs and some related questions. In *Infinite and Finite Sets (to Paul Erd  s on his 60th birthday)*, Colloquia Mathematica Societatis J  nos Bolyai, pages 609–627. J  nos Bolyai Mathematical Society, 1975.
- S. Goldman, R. Schapire, and R. Rivest. Learning binary relations and total orders. *SIAM J. Computing*, 22:46–51, 1993.
- S. Janson. Large deviations for sums of partly dependent random variables. *Random Structures Algorithms*, 24:234–248, 2004.
- R. Jin, S. Wang, and Y. Zhou. Regularized distance metric learning: Theory and algorithm. In *Advances in Neural Information Processing Systems 22*, pages 862–870. 2009.
- V. Koltchinskii and D. Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics*, 30:1–50, 2002.
- M. Ledoux and M. Talagrand. *Probability in Banach spaces: isoperimetry and processes*. Ergebnisse der Mathematik und ihrer Grenzgebiete. Springer-Verlag, 1991.
- C. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics*, volume 141 of *London Mathematical Society Lecture Note Series*, pages 148–188. Cambridge University Press, 1989.
- M. Mohri and A. Rostamizadeh. Rademacher complexity bounds for non-i.i.d. processes. In *Advances in Neural Information Processing Systems 21*, pages 1097–1104. MIT Press, 2009.
- M. Mohri and A. Rostamizadeh. Stability bounds for stationary φ -mixing and β -mixing processes. *Journal of Machine Learning Research*, 11:789–814, 2010.
- L. Ralaivola, M. Szafranski, and G. Stempfel. Chromatic PAC-bayes bounds for non-iid data: Applications to ranking and stationary β -mixing processes. *Journal of Machine Learning Research*, 11:1927–1956, 2010.
- N. Usunier, M.-R. Amini, and P. Gallinari. Generalization error bounds for classifiers trained with interdependent data. In *Advances in Neural Information Processing Systems 18*, pages 1369–1376. MIT Press, Cambridge, MA, 2006.
- L. Valiant. A theory of the learnable. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, pages 436–445, New York, NY, USA, 1984. ACM.
- V. Vizing. On an estimate of the chromatic class of a p-graph. *Diskret. Analiz.*, 3:25–30, 1964.

A Proof of Lemma 5

For each pair of vertices $\{v_i, v_j\}$, define a random variable $E_{i,j} \triangleq \mathbb{1}[\{i, j\} \in E]$. Note that $\deg(v_i) = \sum_{j \neq i} E_{i,j}$ and, via linearity of expectation,

$$\mathbb{E}[\deg(v_i)] = \sum_{j \neq i} \mathbb{E}[E_{i,j}] = (n-1) \cdot \frac{\binom{n}{2}-1}{\binom{n}{2}} = \frac{2m}{n}.$$

Since the expected degree is uniform, let $\mu \triangleq \mathbb{E}[\deg(v_i)]$. Using the union bound, for any $t > 0$, we have that

$$\begin{aligned} \Pr[\Delta(G) \geq t] &\leq \sum_{i=1}^n \Pr[\deg(v_i) \geq t] \\ &= \sum_{i=1}^n \Pr\left[\sum_{j \neq i} E_{i,j} \geq t\right]. \end{aligned}$$

Though the variables $\{E_{i,j}\}_{j \neq i}$ are dependent, it is straightforward to show that they are negatively correlated. Therefore, we can apply multiplicative Chernoff bounds, with $t = \mu(1 + \epsilon)$, to obtain

$$\Pr[\Delta(G) \geq t] \leq \sum_{i=1}^n \exp(-\mu\epsilon^2/3) = n \exp\left(\frac{-2m\epsilon^2}{3n}\right).$$

Setting this equal to δ , then solving for ϵ and t completes the proof.